

“Resource Pooling” for Wireless Networks: Solutions for the Developing World

Junaid Qadir¹, Arjuna Sathiaselan², Liang Wang², Jon Crowcroft²

¹Information Technology University (ITU)-Punjab, Lahore, Pakistan

²Computer Laboratory, University of Cambridge, United Kingdom

Email: junaid.qadir@itu.edu.pk; arjuna.sathiaselan@cl.cam.ac.uk; liang.wang@cl.cam.ac.uk; jon.crowcroft@cl.cam.ac.uk

ABSTRACT

We live in a world in which there is a great disparity between the lives of the rich and the poor. Technology offers great promise in bridging this gap. In particular, wireless technology unfetters developing communities from the constraints of infrastructure providing a great opportunity to leapfrog years of neglect and technological waywardness. In this paper, we highlight the role of resource pooling for wireless networks in the developing world. Resource pooling involves: (i) abstracting a collection of networked resources to behave like a single unified resource pool and (ii) developing mechanisms for shifting load between the various parts of the unified resource pool. The popularity of resource pooling stems from its ability to provide resilience, high utilization, and flexibility at an acceptable cost. We show that “resource pooling”, which is very popular in its various manifestations, is the key unifying principle underlying a diverse number of successful wireless technologies (such as white space networking, community networks, etc.). We discuss various applications of resource pooled wireless technologies and provide a discussion on open issues.

1. INTRODUCTION

It has been shown that the Internet has the potential to foster development and growth, but this potential is thwarted by the fact that billions of people are unable to access the Internet [1]. People not having Internet access—nearly half of whom live below the poverty line in their respective countries—are more likely to be uneducated, poor, and residents of undeveloped rural areas. These availability and affordability gaps are disproportionately impacting people in Africa, Asia, and Latin America and contributing to a “digital divide” between humanity. Bringing the Internet to the remaining billions of people will democratize knowledge, open up new opportunities, and undoubtedly open up avenues for sustained development. The fact that Internet access can play a large role in facilitating development motivates the vision of *Global Access to the Internet for All* (GAIA), currently being formally pursued in the Internet Research Task Force (IRTF).

While traditionally fixed wired broadband connections have been used to provide Internet to the majority of online population in the developed countries, this approach is not attractive for developing countries due to the high cost of wiring and infrastructure. The emergence of *mobile phones* and *wireless technology* allows the people in developing regions an opportunity to bypass the need of fixed-line broadband and personal computers and leapfrog towards an mobile broadband offering [2].

While mobile technology has been one of the biggest success stories of engineering, cellular technology—which is employed by (for-profit) telecom after deploying terrestrial infrastructure such as base stations—is not well suited for provisioning global Internet/ connectivity. Since sufficient user density is needed to justify the cost of installing expensive infrastructure, mobile connectivity is mostly restricted to urban areas. In addition, traditional cellular companies are also hard pressed to scale to the exponentially increasing performance demands (due to exponential rise in the number of mobile users and their bandwidth demands).

While the use of wireless technology in the developing countries has been proposed in a number of papers [3] [4], traditional wireless networks suffer from a number of performance, scalability, and economic viability problems. In this paper, we argue that *resource pooling based wireless technologies*, through its various manifestations (elaborated in Section 2), can provide a solution to the problems associated with traditional wireless. The main idea underlying “resource pooling” is to develop an abstraction of a singled pooled resource that behaves as a single queue even when the resources constitute a disparate collection. Resource pooling techniques characteristically operate in a bottom-up fashion, which has been shown to be an effective operating principle for developmental work [5]. Our work is the first—to the best of our knowledge—that highlights resource pooling as a general underlying architectural principle that can be used to improve the appropriateness of wireless technology for developing regions while also improving its efficiency and performance.

The organization of this paper follows. We describe the various manifestations of resource pooling in Section 2. We describe how the resource pooling principle underpins diverse wireless technologies in Section 3. We discuss the various open issues related to resource pooled wireless technologies in Section 4. We conclude this paper in Section 5.

2. RESOURCE POOLING MANIFESTATIONS

Although resource pooling has many manifestations in wireless technology, these manifestations can be categorized according to seven principal signature attributes. A description of this classification—which is not mutually exclusive, i.e., a technology can embed multiple resource pooling manifestations—follows and is summarized in Table 1.

2.1 Techniques Based On Aggregation/ Concurrency/ Bonding

A common type of a resource pooling technique depends on bonding or aggregating multiple resource together to act

like one bigger more powerful resource. Such a technique also can utilize the aggregated resources concurrently. Examples of this type include the use of multiple paths to communicate with a destination; the use of multi-input multi-output (MIMO) techniques; and bonding multiple channels together, as in Wi-Fi bonding, and 3G bonding.

Table 1: Wireless Resource Pooling Technique Types and Examples

1. Techniques that use <u>aggregation/ concurrency/ bonding</u>:
Multipathing (over multiple wireless interfaces: e.g., 3G/ Wi-Fi) MIMO; Wi-Fi bonding [6]; 3G bonding [7];
2. Techniques that use <u>centralization</u> of resources:
Radio Resource Management (RRM)
Cloud radio access networks (C-RAN) [8]
Mobile Virtual Network Operators (MVNOs)
Wireless spectrum resource pool (WSRP) [9]
3. Techniques that use <u>dynamic allocation of resources</u>
Dynamic Spectrum Access (DSA); White Space Networking
Less than Best Effort (LBE) Service [10]
4. Techniques for <u>load balancing</u> to control congestion:
<i>Load Balancing Over Paths/ Channels/ Peers:</i>
Multihoming; Multichannel wireless
Multipath routing and transport
<i>Load Balancing Over Time:</i>
Low Extra Delay Background Transport (LEDBAT) [10]
Delay Tolerant Networking [11]; Pocket Switched Networks [12]
5. Techniques for <u>Resource sharing/ multi-tenancy/ virt.</u>
Multiplexing, Packet Switching, DSA
Virtualization, Cloud Computing, MVNOs
Virtualization (of BTS; VMs, VNs) [13]
Virtualization of 4G/ 5G RAN; Wi-Fi APs & NICs [14];
6. Techniques based on <u>community sharing/ crowdsourcing</u>:
Wi-Fi offloading [15]
Community Wireless Networks: LCD-Net [16]; Public Access Wi-Fi Service (PAWS) [17];
Community GSM Networks [18];
Fog Computing [19]
7. Techniques for <u>versatile resource usage</u>:
Software Defined Radio (SDR); Network Coding;
TCP Hierarchical Acknowledgement (HACK) [20]

2.2 Techniques Using Centralized Resources

There has been a perennial tension between the distributed and the centralized processing and management paradigms. While distributed processing allows for more robust localized implementations, it is also true that centralized schemes guarantee optimal resource pooling (albeit at the cost of additional communication overhead and lower fault tolerance). It has been shown by Tsitsiklis et al. [21] that resource pooling benefit can be exploited even with a small degree of centralization leading to exponential performance improvements in steady-state scaling of system delay for sufficiently large systems.

2.3 Techniques For Dynamic Allocation of Resources

Decades of experience with Internet has reinforced a general rule of thumb: it is nearly always more preferable to dynamically allocate resources (a concept that complements “resource pooling”) over static resource allocations. Dynamic resource allocation can allow for more efficiency and

flexibility in situations where the demand on scarce resources is not predictable. The *packet switching* technology, which acts as a foundation of the Internet, is essentially a resource pooling technique. Packet switching uses *statistical multiplexing* to achieve resource pooling through dynamic resource allocation by allowing a burst of traffic on a single circuit to use spare capacity on other circuits. More specific for wireless networks, the “*dynamic spectrum access*” (DSA) technique (elaborated in Section 3.1) is another dynamic resource allocation technique. The DSA model tolerates an unlicensed, or the secondary user (SU), in licensed spectrum subject to the provision that the licensed users, or the primary users (PUs), are not interfered with. Another example of dynamic allocation of wireless resource pool is the use of *unlicensed spectrum* to unlock spectrum (which in licensed spectrum would be dedicated for particular licensed applications/ users). The use of unlicensed spectrum has buoyed an increased interest in using Wi-Fi, and similar unlicensed technologies, for bringing Internet access to the developing world (and hundreds of communities are already operating municipal broadband wireless network that provides broadband access to the Internet [22]).

2.4 Techniques Relying On Load Balancing

Load balancing and resource pooling are also closely related. The efficiency and the power of the pooled resource depend on how the load is balanced over the individual resources. Load balancing can be temporal (i.e., the load is balanced between peak and non-peak hours), spatial (i.e., the load is balanced between different locations), or it can be over multiple resources (e.g., the load is balanced over multiple links, as in multihoming, or over multiple paths as is done in multipath transport protocols). This load balancing can be over paths/ channels/ peers or over time (with more illustrative examples provided in Table 1).

2.5 Techniques For Resource Sharing/ Multi-Tenancy

One of the main purpose of the original Internet was to lower the cost of computing by networking the time-sharing systems that existed at many universities. The same aim of resource sharing and multi-tenancy is driving the modern trends of cloud computing and virtualization. Cloud computing relies on pooled resources gathered at centralized datacenters to support the computing requirements of multiple tenants through virtualization. The multi-tenant model of cloud computing is closely related to resource pooling: in a cloud, each tenant is assigned, and reassigned as needed, all the physical and virtual resources it needs per its demand.

2.6 Techniques Based On Crowdsourcing/ Community Sharing

A number of diverse wireless technologies—such as mobile crowdsensing [23]—have been developed that can effectively pool the increasing capabilities of mobile devices to provide business, healthcare, environmental monitoring, transportation, and social networking services. Other wireless resource pooling techniques include crowdsourced efforts such as Wi-Fi offloading aimed at improving the scalability and performance of mobile networks [15]. Community networking techniques are discussed in more detail later in Section 3.2.

2.7 Techniques For Versatile Resource Usage

The Latin word *versatilis* connotes turning, or having the

capable of turning to varied subjects or tasks. One aim of resource pooling is to versatily apply resources in diverse on-demand, and not predefined, configurations. Versatile resource usage pools the same resource for many different purpose. A good example of a versatile resource pooling technique is software defined radio (SDR). SDRs by their versatile nature are radio chameleons that can use software programmability to pool various wireless networking technologies onto the same radio resource—e.g., SDR may run a telephony protocol (e.g. CDMA) at one time and a totally distinct data communication protocol (e.g. Wi-Fi) at another. Other examples of versatile resource pooling include network coding (which pools multiple messages into a smaller number of coded messages by more versatily defining the routing forwarding action) and cross-layer optimization that pools multiple acknowledgments (ACKs) into a single ACK making ACKs more versatile.

3. RESOURCE POOLED WIRELESS (WIRELESS 2.0)

We will next cover a number of important wireless resource pooling techniques that will likely form the basis of “Wireless 2.0”, the next generation of wireless technology that will have better efficiency and also more advanced functionality. In addition to their role in advancing the state of the art, these technologies are also amenable to deployments in developing environments (due to their bottom-up resource-pooled nature).

3.1 Dynamic Spectrum Access

Dynamic spectrum access (DSA) refers to the set of technologies that can enable radios to opportunistically access radio spectrum (that may be licensed but is currently unused) to effectively create a “secondary network” on top of an existing “primary network”. In networks that use DSA, also known as “white space networks”, nodes can increase the efficiency of spectrum utilization through “spectrum pooling” [24] and thereby avoid the problem of “artificial spectrum scarcity” inherent to the command-and-control licensing approach to spectrum that can deny consumers and their devices wireless bandwidth even when no one else is using it. Just like packet switching improved upon circuit switching, DSA can improve the command-and-control spectrum licensing approach by sharing the radio spectrum through statistical multiplexing. However, the increased flexibility due to DSA comes with many technological challenges with the primary challenge being ensuring protection to “primary users” (PUs).

There are two important types of white space networks: TV white space (TVWS) networks and GSM white space networks. TVWS is composed of the VHF and UHF band spectrum that was traditionally earmarked for TV transmissions but was unused. There is an increasing trend globally of opening these TV white spaces (which have very desirable transmission properties) for broadband access. This trend is being accelerated by the switch from analogue to digital television (which takes less spectrum). Broadly speaking, there are four key reasons that ICT for development (ICTD) community is excited about TVWS in developing countries [25]: i) low risk regulation; ii) suitability for rural environments; iii) ample availability of TV spectrum in rural/ underdeveloped areas; and iv) opportunities for entrepreneurship (through which the wireless service market

has been democratized).

With the emergence of low-cost software defined radios (SDRs), along with the availability of open-source software such as OpenBTS¹, a growing recent trend is to utilize *community cellular networking* to support rural area cellular networking using *GSM white spaces*. While building and operating top-down cellular technology, as operated by traditional mobile companies, is very expensive (especially for rural settings), a rural community can support itself with community networking at a very reasonable cost. The trend of using local community networks with unlicensed spectrum promises to ameliorate the issue of digital exclusion of rural communities [26].

3.2 Community Networking

It has been shown by researchers that the adoption of an indigenous bottom-up approach is the key to sustainable development [5]. Community networks follow the bottom-up philosophy and are “run by users, for the users”. Community-based resource pooling can be in the form of an extension of a user’s paid service (as in BT FON², where home users publicly share their home broadband connections in return of credits that may be used to access other user’s APs) or in the form of a cooperative wireless mesh network (e.g., Guifi³, which is an open self-organizing community network that uses unlicensed wireless links and open optical fiber links. The applications of community networking extend to much more than Internet access: e.g., a number of community communication scenarios can be entertained such as local intranets and machine-to-machine networks that provide localized services, such as neighborhood groups and surveillance, high-speed P2P networking, high-speed video conferencing, local TV/radio broadcasting stations. While community networking also has application in urban settings, the biggest draw for community networking is in rural settings (where cellular networks and satellite networks are expensive to operate). If suitably incentivized, crowd-sourced resource pooling has the potential of mitigating the digital exclusion of rural populations by breaking the impasse emerging from the mismatch between the economic (profit-seeking) interests of commercial service providers and the egalitarian goal of GAIA.

3.3 Multihoming with Heterogeneous Networks

In the not too distant past, multihoming was relevant only for critical servers, but with billions of mobile now multihomed with heterogeneous access technologies, multihoming has now become mainstream. There is a lot of ongoing research on supporting heterogeneous multihoming with work aiming to support (i) seamless handover between mobile cellular connections and Wi-Fi [27], (ii) supporting scalable wireless networking through Wi-Fi offloading and 3G on-loading, and (iii) concurrent usage, or bonding, of data and telecom connections for higher performance [28]. A modern multihomed wireless device can exploit its interfaces (such as 3G/ Wi-Fi) to improve its reliability, throughput, latency, and fault tolerance. In recent times, there have been multipath transport level initiatives (such as MP-TCP) that seek to support such multihoming by allowing a client to establish multiple connections to the same destination host over different network adapters. Siri on Apple’s iOS7 uses MP-TCP

¹<http://openbts.org>

²<http://www.btfon.com>

³<http://guifi.net>

to improve its usability by making its service more resilient to network failure. This is managed by opening two connections: the primary TCP connection is created over Wi-Fi while the backup connection—which is used when Wi-Fi becomes unresponsive—works using the cellular connection.

3.4 Mobile Cloud Computing

The mobile cloud computing trend seeks to exploit the great promise of cloud computing and the ubiquity of mobile smartphones. Apart from the benefits of scalability, performance, and efficiency, using cloud computing in developing countries is attractive for another important reason. Users in developing countries can combine cheap battery-powered smartphones with cloud servers to circumvent the power problems due to unreliable intermittently available electrical grid. In addition, the lack of existing infrastructure implies that users in developing countries can directly adopt advanced cloud/ wireless technology without being encumbered by infrastructure conversion costs [29]. Assuming the availability of a decent Internet connection, the “software as a service” model is also attractive in the developing world as a substitute for in-house services like email, file sharing, etc. In addition, a world of opportunities in fields as diverse as mobile health, education, banking, agricultural support exist (through mobile paradigms dubbed as *m-Health*, *m-Education*, *m-Banking*, and *m-Agriculture*) [29]. In related work, there has been a lot of recent efforts to bring the benefits of cloud computing to cellular and wireless networking. In particular the “*network without borders*” (NWoB) architecture presented in [13] envisions a pool of resources (comprising aggregated spectrum and infrastructure resources) that are used and repurposed as needed to support the instantiation and orchestration of multiple virtual wireless networks.

3.5 Virtualized Wireless Networks

Virtualization has transformed the efficiency of the data-center environment by enabling cloud computing and promises to bring a similar revolution to wireless networking. The trend of virtualized wireless networks (VWNs) is fast emerging since it brings the performance, convenience, and cost-efficiency of virtualization to wireless networking [14]. VWNs can facilitate the convergence and interworking of different wireless technologies and enable efficient resource sharing—e.g., VWNs can be used to support multiple virtual network operators (VNOs), using which mobile VNOs can share the cellular infrastructure and thereby lower its capital and operational expenditures (CAPEX and OPEX). In recent times, there have been efforts of virtualizing WLANs, cellular networks, SDN-based wireless networks, as well as cloud-based VWNs and cognitive-radio-based VWNs [14]. There is also the emerging trend of network functions virtualization (NFV), which promises to disrupt the telecom industry by decoupling the functionality of telecom devices from its physical instantiation in dedicated telecom hardware.

3.6 Edge Computing/ Fog Computing

The two main driving forces of edge/ fog computing are: i) the undesirability of moving large amounts of data to the cloud, especially through low-bandwidth wireless links; and ii) the torrential amounts of data that will increasingly characterize internet in the era of IoT. With edge computing/ fog computing, the resource pool at the edge (comprising of the computing, storage, and communication resources of

user’s smartphones/ tablets) is utilized instead of the central resource pool that may be hosted in a cloud datacenter [19]. By exploiting the edge resources and avoiding the reliance on the cloud, we can reap the benefits of better responsiveness and a wider service coverage. The recent trend of device-to-device (D2D) communication—defined in the context of cellular mobile networking as the direct communication between two mobile users (also called UEs or User Equipment) without traversing the base station or the core cellular network—works in a similar vein. D2D communication technology can be a valuable fallback for public safety communications in the case of disasters and emergencies in which some (or all) of the cellular networking infrastructure becomes fails or becomes unavailable. The use of D2D communication in 3G, 4G, and 5G technology is currently being actively pursued.

3.7 Delay-Tolerant Networking (DTN)

Since it is anticipated that many alternative networking technologies will play a part in universal Internet access, the always-connected connectivity model of the Internet does not lend itself as an economically viable solution for GAIA. DTN is a temporal resource pooling technique that can be used to transport traffic over time even when end-to-end connectivity is intermittent. Originally developed as an architecture for the interplanetary Internet [30], delay tolerant networks (DTNs) are now recognized as an alternative solution for building disruption tolerant solutions in systems where the communication links are intermittently available (e.g., as in LEO constellation system). DTN communication paradigm aims to provide a less-than-best-effort service and aggressively widen the possible connectivity options beyond those available currently.

3.8 Information-Centric Networking (ICN)

The basic insights of ICN are (1) the Internet operation should be focused on the exchange of content (without being coupled to the location of the communicating end hosts), and (2) there should be universal caching embedded in the architecture. ICN proposes to replace the the send-receive paradigm of the current Internet (that couples the location/identifier of the communicating machines) will be replaced by a publish-subscribe paradigm (that provides an architectural spatial and time decoupling). ICN promises benefits such as lower response time (due to pervasive caching and nearest-replica routing) and better support for traffic engineering, content distribution, security, and mobility. ICNs can provide many of the benefits of CDNs (such as providing the requested data from close locations) without the drawbacks of CDNs (which are often implemented using proprietary technologies as overlays—and thus cannot utilize the underlying network’s multicast and broadcast capabilities). Similarly, ICN also leverages insights and technology from P2P networking (such as the use of distributed hash tables (DHTs)).

4. DISCUSSION

4.1 Resource Pooling in Time and Space

By employing resource pooling in time and space, we can unshackle the close coupling of the various disparate entities in the original Internet architecture—such as the close coupling of i) the host identifier and location in the case of IP address, and of ii) the content requested and the loca-

tion of a server hosting this content—that hinders a flexible evolution of Internet’s architecture to accommodate new applications such as video and mobility. In recent times, researchers have proposed many architectures that aim to redress such couplings. The DTN architecture aims to provide a decoupling in time, while the ICN architecture aims to provide a decoupling in space. Such decoupling is especially relevant in the context of developing countries and remote areas where end-to-end communication may not be possible and localized opportunistic communication needs to be used more leveraging ideas that follow from spatial and temporal decoupling of communicating nodes. By disrupting the traditional assumption of always-on Internet connectivity, nodes in such challenged networks can communicate using e.g. time-shifted access (in which delay-tolerant asynchronous bulk data is transferred at off-peak hours). In such scenarios, batched transfer (in contrast to always-on end-to-end transfer) can also lead to energy savings since a receiver does not have to continuously check for and process arriving packets. In recent times, researchers have focused their attention on exploiting the resource pooling in time and space through the DTN/ ICN hybrid for enabling global access to the Internet for all (GAIA) [31].

4.2 Where, When, and How to Resource Pool?

4.2.1 Centralized or Distributed Management

While the resource pool incorporates centralization and aggregation of resources, the management of the resource pool can be done in a centralized or distributed fashion. Radio resource management (RRM)—which performs system-level control of radio resources and parameters in technologies such as cellular, wireless, and broadcasting networks—can be performed in both centralized and distributed fashion (although central control is more common). The modern software-defined networking (SDN) trend however is based on the centralization of the network control functionality in a centralized controller.

4.2.2 When (& When Not) To Use Resource Pooling

Throughout this paper, we have provided many examples of how the economics of supporting a large number of users/ applications make the use of resource pooling attractive. This should not be construed to mean that anti-resource-pooling (sometimes called resource fragmentation) is without utility. There is a need of both models (resource pooling/ fragmentation). As an example, in OFDM, the justification for anti-resource-pooling is the ease of signal processing on narrow frequency bands to achieve higher spectral efficiency. Anti-resource-pooling can help avoid ‘the tragedy of the commons’. Example of techniques that rely on resource fragmentation include time division multiplexing (TDM)/ frequency division multiplexing (FDM), orthogonal FDM (OFDM), time division duplex (TDD)/ frequency division duplex (FDD). The main advantage of resource fragmentation is improved service (often at a higher cost) while the main advantage of resource pooling is cost-efficient resource utilization due to the ability to utilize unused resource capacity.

4.2.3 How to Fairly Share The Resource Pool

An important question relating to resource pooling is to decide how to fairly divide the resources (i.e., who gets which

share of the resource pie?). There are various fairness definitions that have been proposed [32]. The naive view that fairness means absolute equality breaks down when we consider that different users have different preferences and utilities for different resources. It may also be preferable to use unequal allocation if all the users individually get a bigger share of the resource pie. Commonly used fairness criteria include *max-min fairness* (maximize the minimum share every one gets), *max utilization* (the resources are maximally utilized), and (*weighted*) *proportional fairness* in which aggregate of (weighted) proportional changes is non-positive for any other allocation). While devising an efficient resource allocation scheme, it is also important to employ an incentive-compatible scheme that encourages users and applications to choose the “best” action using insights from the field of “mechanism design”. On the Internet, since user applications and end hosts are directly in charge of determining their resource share (according to the TCP protocol), end users can selfishly claim more than their fair share and thus cause a tussle between end hosts, the various stakeholders operating the network [33]. An interesting new approach for a fairer, faster Internet [34] is to allow everyone to use as much of the Internet resource pool as they can, but to police and push back those users/ applications that limit the freedom of others when “too much load meets too little capacity” [35].

4.3 Can Resource Pooling Mechanisms Conflict/ Fail?

While it has been shown that resource pooling indeed does help in responding robustly to failures and overloads, we must also note that resource pooling implementations can be complex or even unstable. Some forms of resource pooling—such as dynamic multipath routing—can lead to instability and hysteresis (the phenomenon in which a physical effect lags behind changes in the effect causing it). The approach of resource pooling can fail if the various stakeholders adopt a myopic greedy approach—captured by the economic notion of the “tragedy of the commons” popularized by Hardin in 1968. Like other commons, unlicensed spectrum can also be plagued potentially by the “tragedy of the commons” in which uninhibited access by many users can render spectrum useless. It is also important to manage any “tussle” that results [33] through the interaction of various stakeholders of the Internet (e.g., users and network), and to ensure that the various multi-layer resource pooling techniques combine harmoniously without any conflicts [36].

5. CONCLUSIONS

The main contribution of this paper is that we identify “resource pooling” as a general unifying principle underpinning a wide variety of wireless technologies. We advance the state of the art by demonstrating the generality of the concept of resource pooling in wireless networks. Since the individual wireless technologies based on resource pooling have been developed independently, a plethora of piecemeal solutions have emerged that may unfortunately not work together harmoniously, and may even conflict. A principled unified understanding of resource pooling in wireless technologies can act as the first step towards avoid potentially conflicting wireless resource pooling mechanisms. We also highlight the important role such resource pooling techniques play in the developing world.

6. REFERENCES

- [1] McKinsey, "Offline and falling behind: Barriers to internet adoption," *White Paper*, 2014.
- [2] X. Gao, X. Quan, R. Jain, T. Kawahara, and G. Powell, "Wireless local loop at the bottom of the pyramid," in *IEEE WCNC'2004*, vol. 1, pp. 459–464, IEEE, 2004.
- [3] V. Gunasekaran and F. C. Harmantzis, "Emerging wireless technologies for developing countries," *Technology in Society*, vol. 29, no. 1, pp. 23–42, 2007.
- [4] L. Subramanian, S. Surana, R. Patra, M. Ho, A. Sheth, and E. Brewer, "Rethinking wireless for the developing world," *Hotnets-V*, 2006.
- [5] W. Easterly, "Institutions: Top down or bottom up?," *The American Economic Review*, pp. 95–99, 2008.
- [6] S. H. R. Bukhari, M. H. Rehmani, and S. Siraj, "A survey of channel bonding for wireless networks and guidelines of channel bonding for futuristic cognitive radio sensor networks," *IEEE Communications Surveys and Tutorials*, 2016.
- [7] A. Sathiseelan and J. Crowcroft, "Internet on the move: challenges and solutions," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 1, pp. 51–55, 2013.
- [8] A. Checko, H. L. Christiansen, Y. Yan, L. Scolari, G. Kardaras, M. S. Berger, and L. Dittmann, "Cloud RAN for mobile networks—a technology overview," *Communications Surveys & Tutorials, IEEE*, vol. 17, no. 1, pp. 405–426, 2014.
- [9] M. Yang, Y. Li, D. Jin, L. Su, S. Ma, and L. Zeng, "OpenRAN: a software-defined RAN architecture via virtualization," in *ACM SIGCOMM 2013*, pp. 549–550, 2013.
- [10] S. Shalunov, G. Hazel, J. Iyengar, and M. Kuehlewind, "Low extra delay background transport (LEDBAT)," *draft-ietf-ledbat-congestion-04.txt*, 2010.
- [11] K. Fall, "A delay-tolerant network architecture for challenged internets," in *ACM SIGCOMM*, pp. 27–34, ACM, 2003.
- [12] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Pocket switched networks: Real-world mobility and its consequences for opportunistic forwarding," tech. rep., TR UCAM-CL-TR-617, University of Cambridge, Computer Laboratory, 2005.
- [13] L. Doyle, J. Kibilda, T. K. Forde, and L. Dasilva, "Spectrum without bounds, networks without borders," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 351–365, 2014.
- [14] J. Qadir, N. Ahmed, and N. Ahad, "Building programmable wireless networks: an architectural survey," *EURASIP Journal on Wireless Communications and Networking*, vol. 2014, no. 1, pp. 1–31, 2014.
- [15] K. Lee, J. Lee, Y. Yi, I. Rhee, and S. Chong, "Mobile data offloading: how much can wifi deliver?," in *Proceedings of the 6th International Conference*, p. 26, ACM, 2010.
- [16] A. Sathiseelan and J. Crowcroft, "LCD-net: lowest cost denominator networking," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 2, pp. 52–57, 2013.
- [17] A. Sathiseelan, R. Mortier, M. Goulden, C. Greiffenhagen, M. Radenkovic, J. Crowcroft, and D. McAuley, "A feasibility study of an in-the-wild experimental public access WiFi network," in *Proceedings of the Fifth ACM Symposium on Computing for Development*, pp. 33–42, ACM, 2014.
- [18] S. Hasan, K. Heimerl, K. Harrison, K. Ali, S. Roberts, A. Sahai, and E. Brewer, "GSM whitespaces: An opportunity for rural cellular service," in *IEEE DYSpan'2014*, pp. 271–282, IEEE, 2014.
- [19] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the internet of things," in *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*, pp. 13–16, ACM, 2012.
- [20] L. Salameh, A. Zhushi, M. Handley, K. Jamieson, and B. Karp, "HACK: hierarchical acks for efficient wireless medium utilization," in *Proceedings of the USENIX Annual Technical Conference*, pp. 359–370, 2014.
- [21] J. N. Tsitsiklis, K. Xu, *et al.*, "On the power of (even a little) resource pooling," *Stochastic Systems*, vol. 2, no. 1, pp. 1–66, 2012.
- [22] M. Mandviwalla, A. Jain, J. Fesenmaier, J. Smith, P. Weinberg, and G. Meyers, "Municipal broadband wireless networks," *Communications of the ACM*, vol. 51, no. 2, pp. 72–80, 2008.
- [23] N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. T. Campbell, "A survey of mobile phone sensing," *Communications Magazine, IEEE*, vol. 48, no. 9, pp. 140–150, 2010.
- [24] T. A. Weiss and F. K. Jondral, "Spectrum pooling: an innovative strategy for the enhancement of spectrum efficiency," *Communications Magazine, IEEE*, vol. 42, no. 3, pp. S8–14, 2004.
- [25] E. E. Pietrosemoli and M. Zennaro, *TV White Spaces: A Pragmatic Approach*. Cambridge university press, 2013.
- [26] K. Heimerl, S. Hasan, K. Ali, E. Brewer, and T. Parikh, "Local, sustainable, small-scale cellular networks," in *Proceedings of ICTD'2013*, pp. 2–12, ACM, 2013.
- [27] C. Paasch, G. Detal, F. Duchene, C. Raiciu, and O. Bonaventure, "Exploring mobile/WiFi handover with multipath TCP," in *Proceedings of the 2012 ACM SIGCOMM workshop on Cellular networks: operations, challenges, and future design*, pp. 31–36, ACM, 2012.
- [28] S. Deng, R. Netravali, A. Sivaraman, and H. Balakrishnan, "WiFi, LTE, or both?: Measuring multi-homed wireless internet performance," in *Proceedings of the 2014 Conference on Internet Measurement Conference*, pp. 181–194, ACM, 2014.
- [29] S. Greengard, "Cloud computing and developing nations," *Communications of the ACM*, vol. 53, no. 5, pp. 18–20, 2010.
- [30] S. Burleigh, A. Hooke, L. Torgerson, K. Fall, V. Cerf, B. Durst, K. Scott, and H. Weiss, "Delay-tolerant networking: an approach to interplanetary internet," *Communications Magazine, IEEE*, vol. 41, no. 6, pp. 128–136, 2003.
- [31] D. Trossen, A. Sathiseelan, and J. Ott, "Towards an information centric network architecture for universal internet access," *ACM SIGCOMM Computer Communication Review*, vol. 46, no. 1, pp. 44–49, 2016.
- [32] T. Lan, M. Chiang, *et al.*, "An axiomatic theory of fairness in network resource allocation," *2010 Proceedings IEEE INFOCOM*, pp. 1–9, 2010.
- [33] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, "Tussle in cyberspace: defining tomorrow's Internet," in *ACM SIGCOMM Computer Communication Review*, vol. 32, pp. 347–356, ACM, 2002.
- [34] B. Briscoe, "A fairer, faster internet," *Spectrum, IEEE*, vol. 45, no. 12, pp. 42–47, 2008.
- [35] A. Jacquet, B. Briscoe, and T. Moncaster, "Policing freedom to use the internet resource pool," in *Proceedings of the 2008 ACM CoNEXT Conference*, p. 71, ACM, 2008.
- [36] D. Wischik, M. Handley, and M. B. Braun, "The resource pooling principle," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 5, pp. 47–52, 2008.